

Multi-material interface reconstruction using particles and power diagrams

S. P. Schofield¹, R. V. Garimella¹, M. M. Francois², and R. Loubere³

1. Mathematical Modeling and Analysis (T-7), Los Alamos National Laboratory

2. Continuum Dynamics (CCS-2), Los Alamos National Laboratory

3. CNRS and Toulouse University, France

December 18, 2006

Abstract

We have developed a new, multi-material, piecewise linear interface reconstruction method that correctly locates the position of each material in the cell and matches the required volume fractions independent of the order in which the materials are specified. This is different from other volume tracking PLIC methods in which an improper ordering may result in materials being incorrectly located within the cell. The new method utilizes a particle attraction model to locate the materials and a form of weighted Voronoi diagram, known as a power diagram, to reconstruct the interface. It works on unstructured grids, for an arbitrary number of materials and is naturally extended to three dimensions.

1 Introduction

The effective management and capture of interfaces is essential to accurate and reliable simulation of multi-material and multi-phase flows. Due to their strict conservation of materials, volume-of-fluid (VOF) methods using interface reconstruction are widely used and will be the focus of this article. VOF methods do not explicitly track the interface between materials, but rather advect volume fractions which prescribe the material composition of each cell of the mesh. When the interface between materials is needed, the interface is recreated based on the material volume fraction in the cell and its surrounding cells [4, 6, 19, 20].

A common problem impacting these reconstruction methods is their dependence on a specified material ordering. If more than two materials are present in a cell, the reconstruction may depend on the sequence in which the materials are processed. This is undesirable as

it may improperly locate materials within the cell. In a finite volume implementation, this may result in material being incorrectly fluxed into neighboring cells.

In this article, we demonstrate a method that can reconstruct a multi-material interface with no dependence on material ordering. The method is very general: it works on unstructured grids, accommodates an arbitrary number of materials and extends naturally to three dimensions. The method utilizes a particle attraction model to infer the relative location of the materials in the cell. Using that information, the interface is reconstructed using a weighted Voronoi diagram, known as a power diagram, such that the required volume fractions are matched.

2 Volume-of-fluid methods

The volume of fluid method, originally developed by C. W. Hirt and B. D. Nichols [14], advects the fractional volumes of each fluid in the cell to track materials in an incompressible flow simulation. The volume fraction, f_m , of a material, m , in a cell with volume A_i is defined as

$$f_m = \frac{A_m}{A_i} \quad (1)$$

where A_m is the volume of the material in the cell. Early VOF methods used a simple interface that was defined to be a coordinate axis aligned line within each cell that partitioned the cell into the correct volume fraction. This is often referred to as the simple line interface calculation (SLIC) due to Noh and Woodward [18]. This interface structure was natural when combined with directionally split advection.

D.L. Youngs [22, 23] extended the method to permit the material interface to have an arbitrary orientation within the cell. In Youngs' method, the normal to an interface within a cell is taken to be the negative gradient of the volume fraction function, estimated using the volume fractions of that material in the neighboring cells. A line intersecting the cell with the prescribed normal is located within the cell so that it cuts off the desired volume. This and other methods that allow an arbitrarily oriented linear interface within a cell are referred to as piecewise linear interface calculation (PLIC) methods [19]. In general, the methods are first order; however, there are extensions that make the reconstruction second order using a local optimization [16] or interface smoothing [11, 21].

While PLIC methods work well for two material cells, typical reconstruction of the interfaces in cells containing more than two materials relies on a material ordering. PLIC methods work by first estimating the interface normals for each material based on the volume fractions of that material. Next the interface for the first material in the ordering is constructed in each cell, separating that material from the rest of the cell, leaving a truncated cell. Then, in the remaining portion of each cell, the interface for the second material is constructed leaving a smaller truncated cell. The process repeats until all the materials have been processed. Unfortunately, the result of this method depends on the order in which the materials are processed. This dependency and extensions of Youngs' method to handle




Figure 1: (a) Correct reconstruction and (b),(c), (d) 1st order PLIC based reconstruction using different material orderings. The numbers designate the order in which the materials were processed.

multi-material reconstructions are explained in greater detail in [9]. The effects of material order dependency are shown in Figure 1. As shown in the top right, with the correct material ordering shown, the interface reconstructed by a first order PLIC method, is close to the correct configuration, and would be identical to the correct configuration after application of a smoothing step or using a second order reconstruction method. Using an incorrect ordering results in substantial degradation of the interface as shown in the two bottom reconstructions.

These incorrect reconstructions may adversely impact the material advection in the simulation. If the advection routine is based on fluxing volumes through the edges of the cell, the improper material ordering may result in a material being fluxed early or not at all. For example, if the flow is moving towards the top right in the bottom left reconstruction, the white material will move into the top right cell prematurely. This can lead to a breakup of the interface.

Selecting a global ordering can be problematic as the appropriate ordering for one region of the mesh may be quite wrong for another. To remedy this, there has been some work on deriving the material order. The geometrically derived material priority by S. Mosso and S. Clancy [17] is based on the assumption of a layer structure and works by approximating the local center of mass of each material, then based on the relative locations along a line, it will

select an ordering. A similar approach was developed by D. J. Benson [5]. However, both methods can fail in the presence of a triple point.

3 Particle and power diagram based reconstruction

In order to eliminate the material order dependence in multi-material interface reconstruction, we have developed a novel method that estimates the location of each material in a cell. Based on that information, the method reconstructs an interface that maintains the correct volume fractions. The method is completely general, working on arbitrary polygonal grids with an arbitrary number of materials in each cell. In addition, it can be naturally extended to three dimensions.

Our method consists of three steps as illustrated in Figure 2:

1. Particles representing the materials are distributed in the cells of the mesh around material interfaces and evolved according to a particle attraction and repulsion model.
2. The particles converge to their final locations from which the relative positions of the materials in the cell are inferred.
3. Based on the material positions, the interface is reconstructed using a weighted Voronoi diagram, known as a power diagram, such that it matches the required volume fractions.

In contrast to existing SLIC or PLIC methods, all materials are processed simultaneously and, as such, have no material order dependency. Furthermore, unlike methods such as a triple point tracking method of Choi and Bussmann [8], no assumption of the topology must be made. The reconstruction will automatically give either the appropriate layer structure or triple point configuration.

3.1 Particle model

In the first step of the method, a number of particles representing the materials, are placed in multi-material cells and any pure or mixed neighboring cells. A particle, P_i , has a position, \mathbf{x}_i , velocity $\mathbf{V}_i = \frac{d\mathbf{x}_i}{dt}$ and material $m(i)$ and is constrained to stay within the cell in which it is initially placed.

Taking inspiration from molecular dynamics [1, 13] and smoothed particle hydrodynamics [7, 12], we evolve the particle positions according to “forces” based on the particles’ relative locations and materials. The velocities of the particles are updated through time integration of a set of ordinary differential equations,

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{V}_i \tag{2}$$

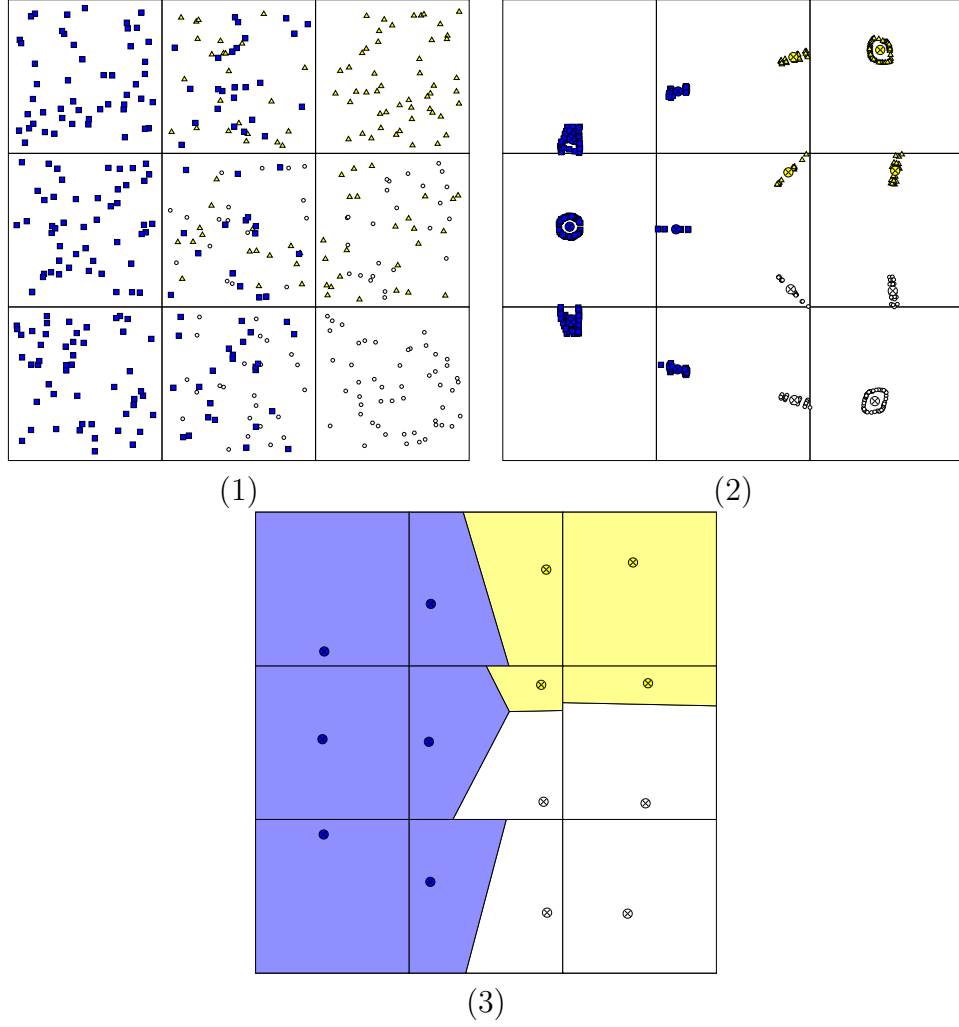


Figure 2: Steps in particle/power diagram based reconstruction: (1) particles are randomly distributed near the interface according to which materials are in each cell. (2) the particle model is run and the aggregate position of the particles is obtained (designated by \otimes) (3) the interface is reconstructed using a power diagram with the generators being the aggregate positions obtained in (2)




Figure 3: Particle attraction and repulsion “forces” used in the model.

$$\mathbf{V}_i = \sum_{j: m(j)=m(i)} \mathbf{V}_{\text{att}}(\mathbf{x}_i, \mathbf{x}_j) + \sum_{j: m(j) \neq m(i)} \mathbf{V}_{\text{rep}}(\mathbf{x}_i, \mathbf{x}_j) \quad (3)$$

where, \mathbf{V}_{att} and \mathbf{V}_{rep} are the prescribed attractive and repulsive “forces” in the direction $\mathbf{x}_j - \mathbf{x}_i$. Particles of the same material attract each other until they are very close, at which point they start to repel each other. Particles of different material repel each other. In our tests, the particles start at random locations within their cell, but they can be initialized using other means such as their relative locations in a cell at a previous time step.

The particle-particle “forces” (plotted in Figure 3) are prescribed as

$$\begin{aligned} \mathbf{V}_{\text{att}}(\mathbf{x}_i, \mathbf{x}_j) &= \begin{cases} -1, & d_{ij} < \delta \\ 1 - 2d_{ij}^4 + d_{ij}^8, & \delta \leq d_{ij} \leq 1.0 \\ 0, & d_{ij} > 1.0 \end{cases} \\ \mathbf{V}_{\text{rep}}(\mathbf{x}_i, \mathbf{x}_j) &= \begin{cases} -(1 - 2d_{ij}^4 + d_{ij}^8), & d_{ij} \leq 1.0 \\ 0, & d_{ij} > 1.0 \end{cases} \end{aligned} \quad (4)$$

where $d_{ij} = \frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{2.5h}$ is the distance between points scaled by an interaction distance, taken to be 2.5 times the characteristic mesh size h , and $\delta = 0.05$.

In a cell, \mathcal{C}_i , the number of particles, $N(\mathcal{C}_i)$ is

$$N(\mathcal{C}_i) = \left\lfloor N_p \times \frac{A_i}{A_0} \right\rfloor \quad (5)$$

where N_p is a prescribed constant (usually around 30), A_i is the area of the cell, A_0 is a reference cell area for the grid, for example on a uniform Cartesian grid, $A_0 = h^2$ where h is the grid spacing and $\lfloor a \rfloor$ is the floor function giving the greatest integer less than or equal to a . Each particle has a designated material type, corresponding to a material present in the cell. Each material that is present in the cell is represented by the same number of particles, $N(\mathcal{C}_i)/N_m^i$, where N_m^i is the number of materials present in the cell. We found

that making the number of particles representing each material proportional to the volume fraction of the material often leads to unsatisfactory results. If the volume fraction is small, the material will be represented by only a few particles, which are not sufficient to provide a reliable estimate of the location of the material within the cell. In addition, we found that for unstructured, general polygonal grids, making the number of particles proportional to the area of the cell was important. Otherwise, the particles tend to cluster in regions of the mesh with a concentration of smaller cells.

Once the particles are distributed, the particle model is run. Since the model prescribes instantaneous velocities and not true forces, the particles may remain in perpetual motion unless the system is forced to “cool”. The velocity of each particle is rescaled at each time step to force the particles to settle into a final configuration. At time $t = k\Delta t$, the kinetic energy of all the particles is

$$KE(t) = \sum_i \frac{1}{2} \|\mathbf{v}_i(t)\|^2 \quad (6)$$

After the first 5 time steps, we force the kinetic energy to decrease as

$$KE^{n+1} \leq \alpha KE^n \quad (7)$$

where $0 < \alpha < 1$. In practice, α is set to be $0.7 - 0.9$. If $KE^{n+1} \geq KE^n$, all the particle velocities are scaled as

$$\mathbf{V}'_i = \sqrt{\alpha \frac{KE^n}{KE^{n+1}}} \mathbf{V}_i. \quad (8)$$

At each timestep, a new Δt is calculated as

$$\Delta t = \frac{0.1}{2\|\mathbf{V}_{max}\|} \quad (9)$$

where $\|\mathbf{V}_{max}\| = \max_i \|\mathbf{V}_i\|$ where \mathbf{V}_i is as defined in Equation 2.

The position is then updated as

$$\mathbf{x}_i^{n+1} = \mathbf{x}_i^n + \Delta t \mathbf{v}_i' \quad (10)$$

If a particle goes outside the cell, it is placed back in the cell by repositioning it to the center of the triangle formed by the old position, the new position, and the center of the cell. If that fails, the particle is kept in its old position. The particles are allowed to evolve for a number of time steps until the average kinetic energy of a particle has dropped below a specified stopping criteria.

3.1.1 Particle clustering

Once the velocity of the particles has dropped to a sufficiently low threshold, the positions of the materials in the cell need to be derived from the final location of the particles. The material location derived from the position of the particles representing the material should




Figure 4: Converged particle locations for a thin filament. The average positions for the particles of each material in a cell are designated by the \otimes markers.

approximate the center of mass of that material within the cell. However, the particles of a material may form multiple groups. A clustering algorithm is needed to detect the multiple clusters and utilize that information to capture the subcell structure. A naive averaging of the particle positions for each material can yield reasonable results if each material in the cell is accurately described by a single convex polygon, but it will not detect the presence of multiple clusters of particles. This is shown in Figure 4. The particles reflect the thin filament structure present, however, the naive averaging used does not adequately reflect that structure. Current research is directed at finding suitable clustering algorithms and appropriate means to divide the volume fractions for a material between multiple clusters.

3.1.2 Convergence

The particle model exhibits rapid convergence to the particle clusters, usually requiring under 20 time steps to converge to approximately the final positions. To demonstrate the convergence with the timesteps, we define two error norms, per material, as

$$E_{avg}(m) = \frac{1}{N} \frac{1}{h} \sum_{i=1}^N \|\mathbf{x}^n(\mathcal{C}_i, m) - \mathbf{x}^\infty(\mathcal{C}_i, m)\| \quad (11)$$

and

$$E_{max}(m) = \frac{1}{h} \max_{i=1 \dots N} \|\mathbf{x}^n(\mathcal{C}_i, m) - \mathbf{x}^\infty(\mathcal{C}_i, m)\| \quad (12)$$

where $\mathbf{x}^n(\mathcal{C}_i, m)$ is the current particle cluster for material m in mesh cell \mathcal{C}_i and $\mathbf{x}^\infty(\mathcal{C}_i, m)$ is the “final” converged particle cluster position. The average is taken over all of the N




Figure 5: The three material configuration used for the particle convergence tests.




Figure 6: The (a) average (E_{avg}) and (b) maximum (E_{max}) error in the location of the particle clusters as a function of the time steps of the particle model. The particle clusters converge to within 5 percent of the mesh spacing within 20 timesteps. Here $N_p = 30$ and $\alpha = 0.75$. For the first 4 time steps, the time step, Δt is held fixed. It is adjusted each subsequent time step.

cells in the mesh. The test case used was that for a filament cutting across the mesh as shown in Figure 5. The convergence history of the particle clusters is shown in Figure 6. This demonstrates that even in the presence of multiple three material cells, the aggregate particle locations converge using only a small number of timesteps.

3.1.3 Initial condition sensitivity

Although the particles are initially randomly distributed, they converge to the same location reliably. We are not concerned with the location of any one particle, but rather the location of the particle cluster, from which the relative location of the material in the cell is determined. A convergence test was run to determine the variability of the particle cluster locations due to the random initial distribution. The configuration chosen was again the test shown in Figure 5. The statistics were generated over $N_s = 100$ runs.

We define two measures of the variability in the particle clusters

$$\sigma_{max} = \frac{1}{h} \max_{\mathcal{C}_i, i=1\dots N} \max_{m \in M(\mathcal{C}_i)} \sigma(\mathcal{C}_i, m), \quad (13)$$

and

$$\sigma_{avg} = \frac{1}{h} \frac{1}{N_s} \sum_{i=1}^N \frac{1}{M(\mathcal{C}_i)} \sum_{m \in M(\mathcal{C}_i)} \sigma(\mathcal{C}_i, m). \quad (14)$$

where $\sigma(\mathcal{C}_i, m)$ is the standard deviation of the particle clusters over the N_s runs for material m in cell \mathcal{C}_i containing the materials $M(\mathcal{C}_i)$. The local standard deviation, $\sigma(\mathcal{C}_i, m)$, is defined as

$$\sigma(\mathcal{C}_i, m) = \left(\frac{1}{N_s} \sum_{k=1}^{N_s} \|\mathbf{x}_k(\mathcal{C}_i, m) - \bar{\mathbf{x}}(\mathcal{C}_i, m)\|^2 \right)^{\frac{1}{2}}. \quad (15)$$

The mean location $\bar{\mathbf{x}}(\mathcal{C}_i, m)$ is taken as the average over all the runs for all parameters considered, that is

$$\bar{\mathbf{x}}(\mathcal{C}_i, m) = \frac{1}{N_s} \sum_{k=1}^{N_s} \mathbf{x}_k^\infty(\mathcal{C}_i, m) \quad (16)$$

where $\mathbf{x}_n^\infty(\mathcal{C}_i, m)$ is the converged aggregate particle location determined in simulation k for $k = 1, \dots, N_s$.

For the filament test case, the variability is shown in Figure 7. The standard deviation per cell averaged over the mesh, represented by σ_{avg} , is around two percent of a cell size, which is quite good. The worst case variability, represented by σ_{max} , is around twenty five percent of the mesh size. This is usually found in cells on the boundary of the computational domain where particles are underrepresented due to the lack of special treatment at the boundary. In general, the average performance improves with the number of particles and with a slower cooling rate, α . However, for a substantial range, the differences are quite small. In our simulations, $N_p = 30$ and $\alpha = 0.75$ provide reasonable results without excessive computational cost.

3.2 Power diagrams

Once the particle groups are found, the interface within the cell is constructed using a power diagram with the point generators being the points determined by the clustering algorithm. A Power diagram or Laguerre diagram [3, 15] is a generalization of a Voronoi diagram generated from a set of points, S , each with an associated radius or weight. The Laguerre distance from a point $\mathbf{x} \in \mathbb{R}^n$ to a point mass, $s_i \in S$ with $s_i = (\mathbf{x}_i, w_i)$ is defined as

$$d_L^2(\mathbf{x}, s_i) = d^2(\mathbf{x}, \mathbf{x}_i) - w_i \quad (17)$$

where $d^2(\mathbf{x}, \mathbf{x}_i) = \sum_{i=1}^n (x - x_i)^2$ is the usual Euclidean distance. If w_i is replaced with w_i^2 in Equation 17, the resulting distance is called the power of the point \mathbf{x} with respect to \mathbf{x}_i .

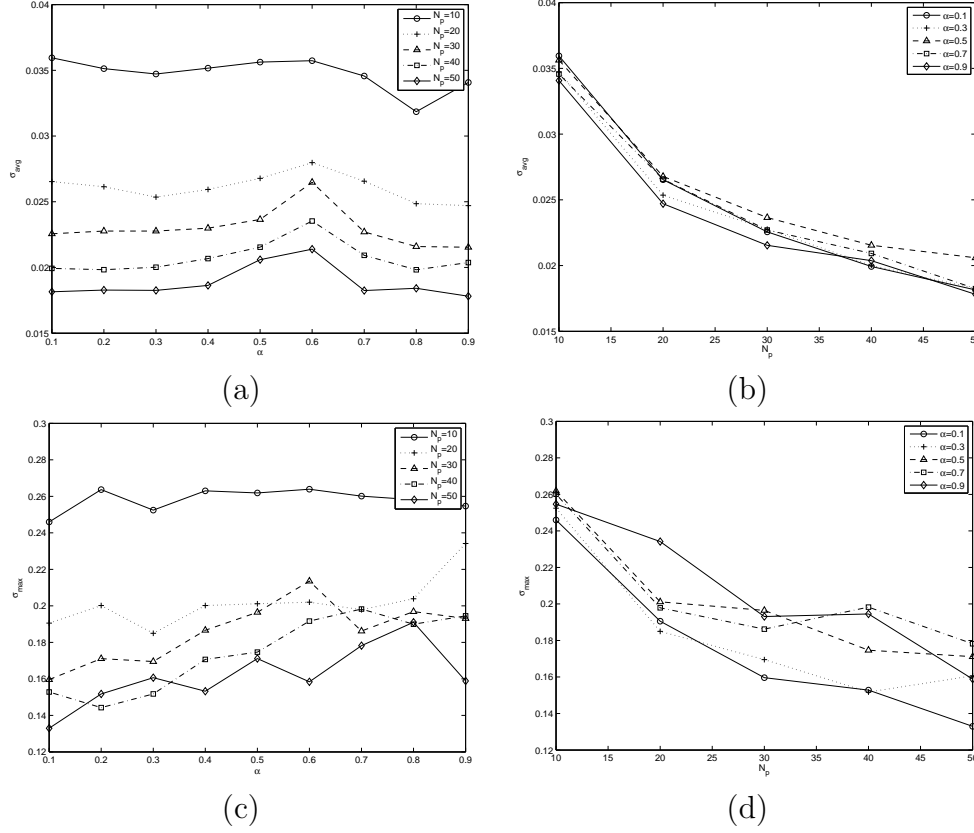


Figure 7: Variability in the final location of the particle clusters depending on the initial conditions. The average standard deviation is within a few percent of the mesh spacing. Larger numbers of particles and values of $\alpha > 0.7$ provide more reliable results.




Figure 8: Power diagram with four generators (V1,V2,V3,and V4) and their weight circles. The generators V1 and V2 do not lie within their corresponding cells.

This may be interpreted as the distance from the point \mathbf{x} to a point on the circle centered at \mathbf{x}_i with radius w_i along its tangent line going through \mathbf{x} .

Each cell in the power diagram is the set of points

$$cell(s_i) = \{\mathbf{x} \in \mathbb{R}^n | d_L^2(\mathbf{x}, s_i) < d_L^2(\mathbf{x}, s_j) \forall s_j \in S, s_j \neq s_i\} \quad (18)$$

The weight associated with a point generator can be interpreted as the square of the radius of a circle centered at that point. Clearly, if all point masses have equal weight (or radius), the power diagram reduces to the usual Voronoi diagram. An example power diagram is shown in Figure 8. Unlike a Voronoi diagram, a point in the point set generating the diagram does not necessarily lie in the cell to which it corresponds. Furthermore, the cell corresponding to a given point mass, may be trivial, that is

$$cell(s_i) = \emptyset \quad (19)$$

In practice, this is not a problem. This is discussed further in Section 3.2.1.

The power bisector (a chordale in Aurenhammer's terminology [3]) between two points $s_i = (\mathbf{x}_i, w_i)$ and $s_j = (\mathbf{x}_j, w_j)$ is the line perpendicular to the segment connecting the points \mathbf{x}_i and \mathbf{x}_j and is located by finding a point, \mathbf{x}_0 on that segment such that $d_L^2(\mathbf{x}_0, s_i) = d_L^2(\mathbf{x}_0, s_j)$.

The power diagram can be constructed in a number of ways. A power diagram may be created through a randomized, incremental algorithm [10], similar to the incremental construction of a Delaunay triangulation. A Voronoi diagram of the point generators may also be efficiently converted into a power diagram [2]. A simple algorithm that intersects all of the mutual power bisectors has $\mathcal{O}(n^2)$ asymptotic complexity in the number of generator points. For our applications, the number of materials in a cell, corresponding to the maximum possibly number of power diagram cells, is small (typically 5 or less), so the asymptotic complexity of the construction algorithm is not a problem. In addition, it has proven to be robust in finite precision arithmetic.

3.2.1 Degeneracies

The reconstructed interface will be the power diagram clipped to the mesh cell. Each of these subcells must match a specified volume. Since power diagrams may have empty cells, it needs to be demonstrated that is not a problem for their use in interface reconstruction. H. Imai [15] provides a useful lemma that provides a sufficient condition for the power diagram cell of a point to be non-trivial:

Lemma (Imai, et al): Given the power diagram for a finite set of point masses, $S = \{s_1, \dots, s_n\}$, $cell(s_i)$ is non-trivial if s_i lies on a corner of the convex hull of S .

If only three materials are present (i.e. S consists of only three point masses), then the point mass corresponding to each material must necessarily be a corner of the convex hull assuming the three points are not collinear. This ensures that each cell, for three material cases, will always be non-trivial for all choices of weights. If the points are collinear, then the cell is partitioned by two parallel lines which can obviously be made to cut off the appropriate volume fractions.

Furthermore, a Voronoi diagram for a set of points will have a non-trivial cell for each point. That implies for every arrangement of (non-coincident) points, there exists a power diagram with all cells having nonzero area, although it may not have the desired the volume fractions.

It remains to be shown that given a set of k points in a convex polygon and the corresponding volume fractions, there is a unique power diagram such that the polygon is subdivided into k subpolygons, each with the appropriate area. While we believe a unique power diagram exists for 4 or more materials in a cell, we have only been able to demonstrate it for the case of 3 or fewer materials.

Also regarding uniqueness, since the power diagram is dictated by the position of the power bisectors and those positions are specified by the relative weights between the point masses, there is a continuous set of weights that will give the same power diagram. Specifically,

Lemma: If P is the power diagram corresponding to the set of point masses (\mathbf{x}_i, ω_i) then for all $\delta > 0$, the power diagram corresponding to the set of point masses $(\mathbf{x}_i, \omega_i + \delta)$ is the same power diagram as P .

Proof: The location of the bisector between two point masses (\mathbf{x}_i, ω_i) and (\mathbf{x}_j, ω_j) is given

by

$$d(\mathbf{x}, \mathbf{x}_i) - \omega_i = d(\mathbf{x}, \mathbf{x}_j) - \omega_j \quad (20)$$

The normal to the bisector always points in the direction from point x^i to x^j . For the power diagrams to be equivalent, it is then necessary and sufficient for all of the bisectors to intersect the line between the points at the same location. Clearly, the point masses $(\mathbf{x}_i, \omega_i + \delta)$ and $(\mathbf{x}_j, \omega_j + \delta)$ have the same bisector as (\mathbf{x}_i, ω_i) and (\mathbf{x}_j, ω_j) .

Lemma: Given a set consisting of 1 to 3 volume fractions, $0 < f_i < 1$, $\sum_i f_i = 1$ and point generators and a bounding, convex polygon, if a power diagram satisfies these volume fractions, then it is unique.

Proof: First consider the case of two points. There is one bisector between the points and its normal is fixed. That leaves one degree of freedom, the position of the bisector along the line between the two points. However, we must match a given volume fraction. This fixes the single degree of freedom. The area behind the intersection line is a monotone function of the position of the bisector guaranteeing a single solution.

In the case of three points, there are now three degrees of freedom (position along the three pointwise lines). If the three points are collinear, then each bisector will partition the polygon into two convex subpolygons, and it is equivalent to repeated application of the two point cases.

Assume the points are not collinear. A property of power diagrams is that the bisectors for any three generators must intersect at a point, i.e. each cell can only have at most two edges arising from the power diagram. Again, the area of a convex subpolygon is a strictly monotone function of the location of the bisector along the line connecting the two points as long as the bisector lies within the bounding polygon. This eliminates the possibility of multiple solutions. Once the two bisectors are fixed to match the volume fraction, the third must automatically follow.

3.2.2 Matching Volume Fractions

The volume fractions can be matched by adjusting the weights of each point generator and checking the area of each cell once it has been clipped to the bounding polygon in which it is contained. This requires the solution of a set of non-linear equations

$$A_m(\omega_1, \dots, \omega_{N_m}) = A_i f_m, \quad m = 1 \dots N_m \quad (21)$$

where $A_m(\omega_1, \dots, \omega_{N_m})$ is the area of the power diagram corresponding to material m after it has been clipped by the bounding polygon with area A_i . f_m is the volume fraction for material m . The constraint

$$\sum_m^{N_m} A_m(\omega_1, \dots, \omega_n) = A_i \quad (22)$$

reduces the number of equations to $N_m - 1$. Specifically, this is done by forcing one of the weights to be a specified value. This enforces a unique set of weights for the desired power diagram.

A Newton procedure with a finite difference derived Jacobian is used to solve Equations (21) and (22). Some caution is required, since the area of each cell is bounded above and below, that is

$$0 \leq A_m(\omega_1, \dots, \omega_n) \leq A_i \quad m = 1, \dots, N_m \quad (23)$$

For extreme values of the weights, some or all of the power diagram cells will be outside of the mesh cell and as such have zero area once clipped to the mesh cell. Furthermore, the A_m will be flat (that is they have a zero gradient) making the Newton procedure fail. As a result, the Newton procedure needs to adjust for overshoots to make sure it does not end up in this region. This is simply done by reducing the size of the Newton step at each iteration if it exceeds those bounds. For the initial guess, we use equal weights for all the point generators. Since the points are all located within the bounding polygon, the areas of each power diagram cell for this initial guess are non-zero. We found the procedure to be robust and efficient, typically converging to tolerance with 3 iterations.

3.3 Accuracy of reconstruction

The use of the power diagram to reconstruct the interface is based on the assumption that we have been able to obtain an approximation to the location of each material in the cell. Indeed, if the point generators for the power diagram are the exact centers of mass of the material subcells, the reconstruction is quite good. For comparison, Figure 9 shows the reconstruction of a four material test case (due to D. S. Bailey, *private communication*) using power diagrams constructed using the true material centers of mass as generators. As is seen in Figure 9, the power diagram based reconstruction is quite good. Of particular note, the location of the materials is preserved in the reconstruction at triple points as well as in layer structures. However, the particle model currently only provides a first-order approximation to the centers of mass. Therefore, the overall method is first order.

V. Dyadechko and M. Shashkov developed a set of three material test cases to evaluate the convergence of multi-material reconstruction methods [9]. The cases consist of three configurations: no junction between the three materials, a T-junction, and a Y-junction. The convergence of the method is evaluated as the radius of curvature of the interface relative to the grid size becomes large. For each test, the true center of mass for each material is used as the generator for the corresponding power diagram cell. The test cases and resulting reconstructions are shown for small and very large radius of curvature in Figures 10 and 11. The power diagram is unable to reproduce the T-junction and, due to the close proximity of the three centers of mass, incorrectly reconstructs the no junction case.

Dyadechko and Shashkov also defined three error norms for interface reconstruction. Here Ω_m is the original (true) region in the cell corresponding to material m , $m = 1 \dots M$, and

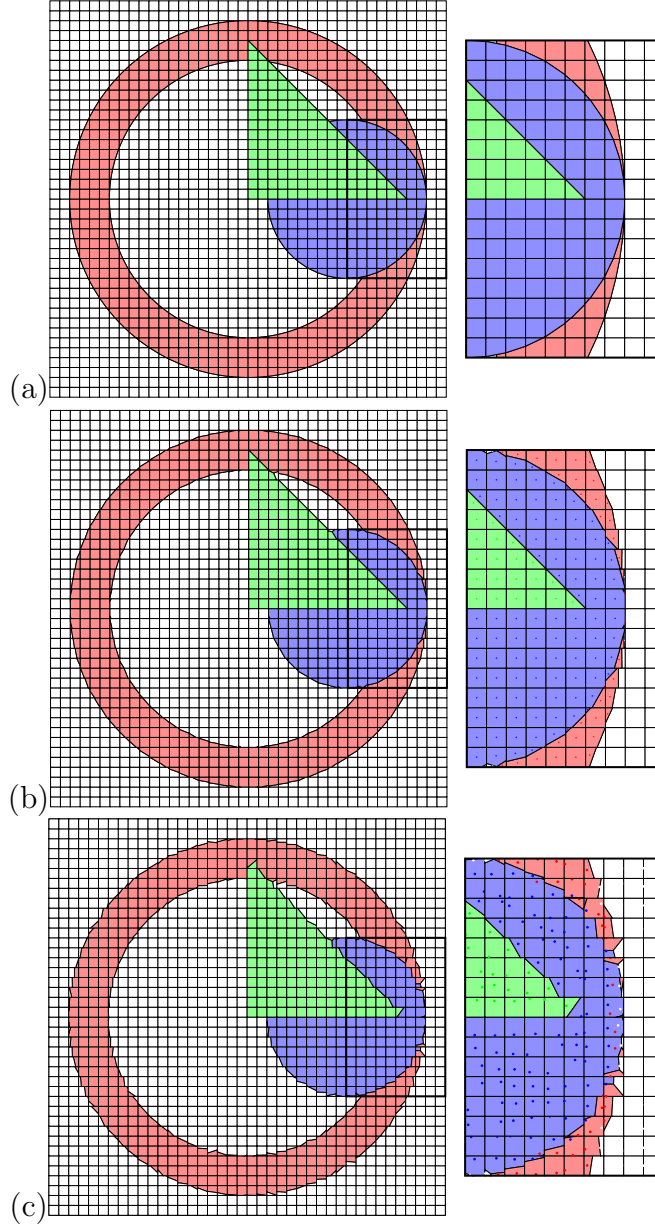


Figure 9: (a) Moment-of-fluid reconstruction (b) Power diagram reconstruction using the material centers of mass as point generators. (c) Power diagram reconstruction using the particle model derived point generators.

Ω'_m is the corresponding region in the reconstruction. Due to matching volume fractions,

$$|\Omega_m| = \int_{\Omega_m} dA = \int_{\Omega'_m} dA \quad (24)$$

The three error norms are:

1. The cumulative defect of the first moment

$$\Delta M_1 = \left(\sum_{m=1}^M \|\mathbf{M}_1(\Omega_m) - \mathbf{M}_1(\Omega'_m)\|^2 \right)^{\frac{1}{2}} \quad (25)$$

where $\mathbf{M}_1(\Omega_m)$ is the first moment of the set Ω_m defined as

$$\mathbf{M}_1(\Omega_m) = \int_{\Omega_m} \mathbf{x} d\mathbf{x} \in \mathbb{R}^2 \quad (26)$$

2. The cumulative area of the symmetric difference

$$\Delta \Omega = \left(\sum_{m=1}^M |(\Omega'_m \cup \Omega_m) \setminus (\Omega'_m \cap \Omega_m)|^2 \right)^{\frac{1}{2}} \quad (27)$$

3. The Hausdorff distance between the boundaries of the regions

$$\Delta \Gamma = \max_{m=1, \dots, M} \max \left\{ \max_{\mathbf{x} \in \partial \Omega'_m} \min_{\mathbf{y} \in \partial \Omega_m} \|\mathbf{x} - \mathbf{y}\|, \max_{\mathbf{x} \in \partial \Omega_m} \min_{\mathbf{y} \in \partial \Omega'_m} \|\mathbf{x} - \mathbf{y}\| \right\} \quad (28)$$

Figures 12, 13 and 14 show the reconstruction error norms as the radius of convergence becomes large. Since the power diagram is unable to exactly reproduce any of the asymptotic configurations, all of the error norms do not continue to decrease with decreasing curvature but rather approach the value of the error associated with the reconstruction of the asymptotic configuration. Note the angles between the interfaces in the power diagram reconstruction of the Y-junction are slightly different than in the actual configuration.

3.4 Stability of reconstruction

The topology prescribed by a Voronoi diagram may change with slight perturbations in the point generators, i.e. cells that did not previously share a common edge may do so after a perturbation. As noted before, the accuracy of the location of a center of mass derived from the particle mode for each material is first order. Therefore, it is possible for small perturbations in the computed centers of mass to lead to the wrong topology for the interface in the cell. In other words, we may obtain a triple point when we expect a layer structure. Of course, mesh refinement can improve the situation.

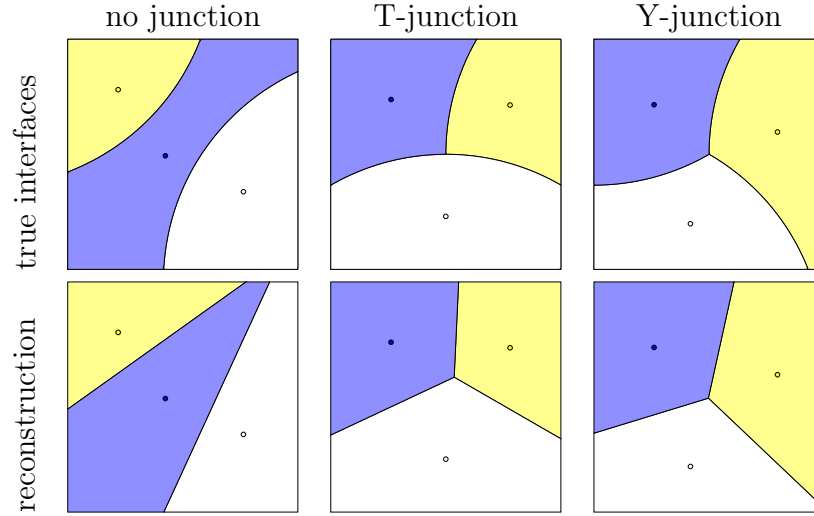


Figure 10: The true material interfaces and their centers of mass are shown in the top row with the power diagram reconstruction using the centers of mass in the bottom row. Here the curvature is set to $\frac{R}{h} = 1$.

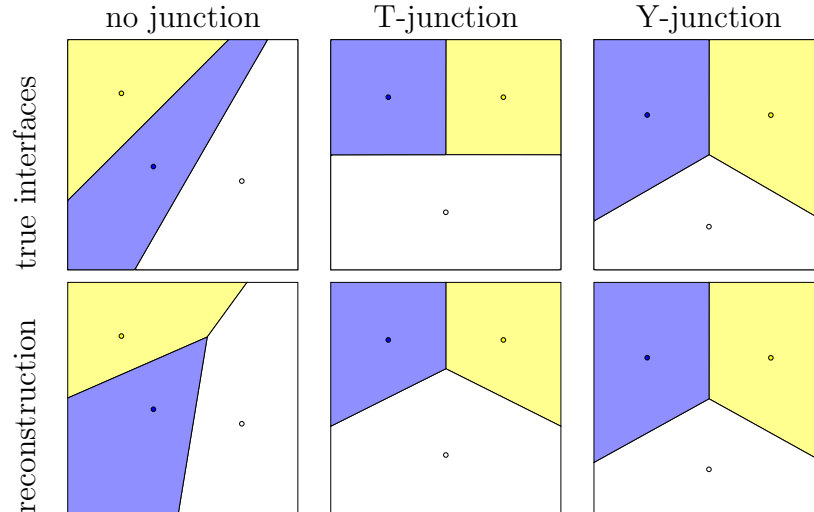




Figure 11: The asymptotic convergence as $\frac{R}{h} \rightarrow \infty$. The power diagrams are unable to recreate the interfaces exactly. Here $\frac{R}{h} = 256$.




$\log_2(R/h)$	ΔM_1 error		
	no junction	T-junction	Y-junction
0	4.21e-02	5.81e-03	2.24e-03
1	5.77e-02	6.69e-03	1.52e-03
2	6.08e-02	7.32e-03	8.96e-04
3	6.09e-02	7.78e-03	6.14e-04
4	6.08e-02	7.90e-03	5.16e-04
5	6.07e-02	7.96e-03	4.89e-04
6	6.06e-02	7.99e-03	4.81e-04
7	6.06e-02	8.01e-03	4.80e-04
8	6.06e-02	8.02e-03	4.79e-04

Figure 12: Cumulative defect of the first moment ΔM_1 as a function of interface curvature.



$\log_2(R/h)$	$\Delta\Omega$ error		
	no junction	T-junction	Y-junction
0	1.42e-01	5.18e-02	3.78e-02
1	1.74e-01	6.34e-02	1.66e-02
2	1.87e-01	6.98e-02	8.47e-03
3	1.90e-01	7.32e-02	5.21e-03
4	1.92e-01	7.49e-02	4.20e-03
5	1.92e-01	7.57e-02	3.88e-03
6	1.93e-01	7.61e-02	3.76e-03
7	1.93e-01	7.64e-02	3.71e-03
8	1.93e-01	7.65e-02	3.70e-03

Figure 13: Cumulative area of the symmetric difference between the true and power diagram reconstructed interfaces ($\Delta\Omega$) as a function of interface curvature.



$\log_2(R/h)$	$\Delta \Gamma$ error		
	no junction	T-junction	Y-junction
0	1.98e-01	9.27e-02	7.16e-02
1	2.66e-01	1.06e-01	4.89e-02
2	3.18e-01	1.16e-01	2.37e-02
3	3.36e-01	1.20e-01	1.41e-02
4	3.45e-01	1.23e-01	9.79e-03
5	3.50e-01	1.24e-01	7.76e-03
6	3.52e-01	1.24e-01	6.78e-03
7	3.53e-01	1.25e-01	6.29e-03
8	3.54e-01	1.25e-01	6.05e-03

Figure 14: Maximum distance between the true and reconstructed interfaces ($\Delta \Gamma$) as a function of interface curvature.




Figure 15: Stability of the power diagram based reconstruction. The true material subcells are shown in dotted lines, the power diagram reconstruction is shown in solid lines.

In Figure 15, the power diagram based reconstruction is shown for a series of perturbed point generators. The volume fractions for each of the material are the same in each reconstruction. For the initial, symmetric case, the top and bottom material do not share an interface. However, as the points are perturbed, they do share an interface. In addition, the top material loses its presence on the right edge of the cell and the middle material moves to have a presence on the top edge. With an edge based flux, this can lead to the incorrect materials being fluxed into the adjacent cells. While this is problematic, there is a continuous dependence of the power diagram with regard to the approximate material locations. The small perturbations in the top row of Figure 15 do not lead to a different topology. Ultimately, the problem is that the material interface structure is not adequately resolved on a coarse mesh. The reconstructions, even with the perturbed point generators, are not unreasonable.

4 Numerical experiments

To assess the overall performance of the method, a number of numerical experiments were performed. Figure 16 shows the convergence of the particles from their initial random distribution. For illustration purposes, the power diagram based interface reconstruction that would result from each particle configuration is shown. The interface is not typically com-

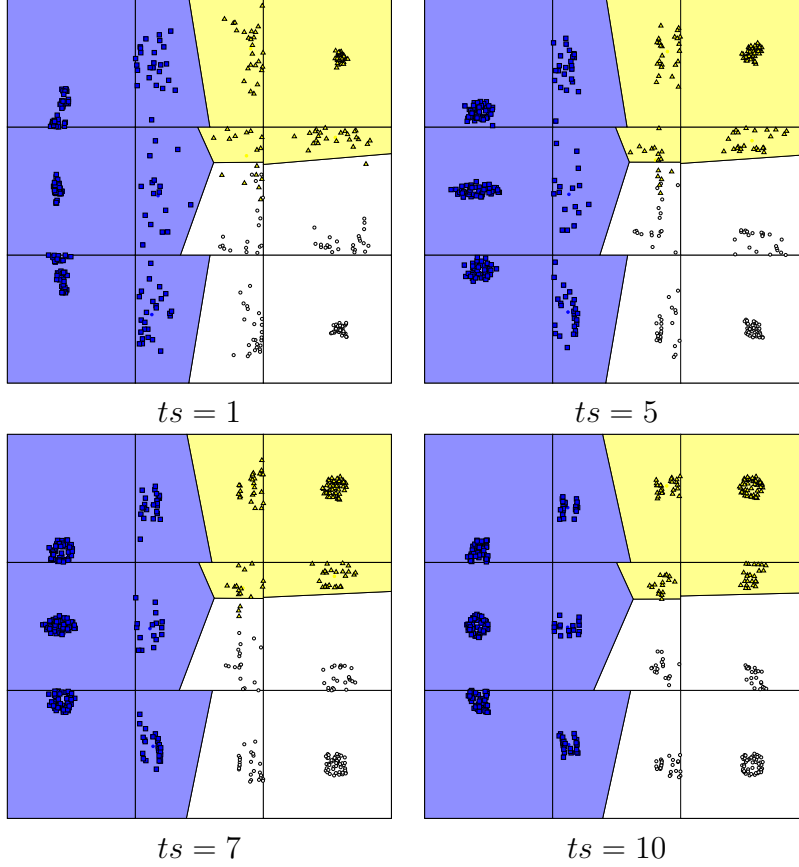


Figure 16: Evolution of the particles through 10 time steps. The power diagram interface reconstruction that would result from each configuration is shown

puted until that particles have converged to their final locations.

In Figure 17, a four material test case is shown. On the structured grid, the reconstruction respects the symmetry in the problem. However, for the unstructured grid, it does not reproduce a straight line, indicating that the reconstruction cannot be second order [16]. Figure 18 shows the reconstruction of a filament type structure that is not aligned with the grid. The wide filament is preserved with all the three material cells showing the proper material positions. The particle/power diagram reconstruction does not reverse the location of the materials relative to the filament as does Youngs' reconstruction. The reconstruction in the cells at the top and bottom of the grid could be improved with a more appropriate boundary treatment. The narrow filament in Figure 19 exhibits breakup which is not uncommon for interface reconstruction methods when applied to fine structures.




Figure 17: Four material example on a (a) structured and (b) unstructured grid. The inset in the top right corner shows the four material cell in the center of the grid with the particles in their final locations.




Figure 18: Example filament structure. (a) particle/power diagram reconstruction (b) Youngs' reconstruction with incorrect material order. Notice the presence of the white material on the right hand side of the filament in Youngs' reconstruction.




Figure 19: Narrow filament structure exhibiting interface breakup. (a) particle convergence (b) resulting power diagram based reconstruction

5 Conclusions

We have developed a first order method for material order independent reconstruction of multi-material cells. It is based on the utilization of a particle model attraction model where particles representing the different materials evolve and converge to locations representing the approximate position of each material within a cell. From the averaged particle positions, an interface is reconstructed using a power diagram, that once clipped to the mesh cell, matches the desired volume fractions.

The performance of the particle model is quite good, in that the average particle position converges to within a few percent of its converged position in under 20 timesteps. Furthermore, the particles do not show much sensitivity to the random initial conditions used.

The power diagram based interface reconstruction maintains the relative location of the materials within the cell and may be useful in other methods utilizing approximate material location information.

As two material reconstruction algorithms are quite good and well established, this method would be most applicable for reconstruction of only cells containing more than two materials. For those cells, the particles would be distributed in a region around the cells to be reconstructed. The converged locations would then be used to reconstruct the cells.

References

- [1] M. P. Allen and D. J. Tildesley, *Computer simulation of liquids*, Oxford University Press, New York, 1987.
- [2] F. Anton and D. Mioc, *On the conversion of ordinary Voronoi diagrams into Laguerre diagrams*, The Eleventh Canadian Conference on Computational Geometry (Vancouver, Canada), August 1999, pp. 150–153.
- [3] F. Aurenhammer, *Power diagrams: properties, algorithms and applications*, SIAM J. Computing **16** (1987), no. 1, 78–96.
- [4] D. J. Benson, *Computational methods in Lagrangian and Eulerian hydrocodes*, Computer Methods in Applied Mechanics and Engineering **99** (1992), 235–394.
- [5] ———, *Eulerian finite element methods for the micromechanics of heterogeneous materials: Dynamic prioritization of material interfaces*, Computer Methods in Applied Mechanics and Engineering **151** (1998), 343–360.
- [6] ———, *Volume of fluid interface reconstruction methods for multi-material problems*, Applied Mechanics Review **55** (2002), no. 2, 151–165.
- [7] W. Benz, *Smooth particle hydrodynamics: a review*, The Numerical modelling of nonlinear stellar pulsations : problems and prospect, NATO ASI series. Series C, Mathematical and physical sciences ; no. 302., Kluwer Academic Publishers, 1989, pp. 269–287.
- [8] B. Y. Choi and M. Bussmann, *A piecewise linear approach to volume tracking a triple point*, International Journal for Numerical Methods in Fluids (in press).
- [9] V. Dyadechko and M. Shashkov, *Multi-material interface reconstruction from the moment data*, Tech. Report LA-UR-06-5846, Los Alamos National Laboratory, Los Alamos, NM, 2006.
- [10] H. Edelsbrunner and N. R. Shah, *Incremental topological flipping works for regular triangulations*, Algorithmica **15** (1996), 223–241.
- [11] R. V. Garimella, V. Dyadechko, B. K. Swartz, and Mikhail J. Shaskov, *Interface reconstruction in multi-fluid, multi-phase flow simulations*, Proceedings of the 14th International Meshing Roundtable (San Diego, CA), Springer, September 2005, pp. 19–32.
- [12] R. A. Gingold and J. J. Monaghan, *Smoothed particle hydrodynamics - theory and application to non-spherical stars*, Monthly Notices Royal Astronomical Society **181** (1977), 375–389.
- [13] J. M. Haile, *Molecular dynamics simulation : elementary methods*, Wiley-Interscience, New York, 1992.
- [14] C. W. Hirt and B. D. Nichols, *Volume of fluid (vof) method for the dynamics of free boundaries*, Journal of Computational Physics **39** (1981), 201–225.

- [15] H. Imai, M. Iri, and K. Murota, *Voronoi diagram in the Laguerre geometry and its applications*, SIAM J. Computing **14** (1985), no. 1, 93–105.
- [16] J. E. Pilliod Jr. and E. G. Puckett, *Second-order accurate volume-of-fluid algorithms for tracking material interfaces*, Journal of Computational Physics **199** (2004), 465–502.
- [17] S. Mosso and S. Clancy, *A geometrically derived priority system for young’s interface reconstruction*, Tech. Report LA-CP-95-0081, Los Alamos National Laboratory, Los Alamos, NM, 1995.
- [18] W. F. Noh and P. Woodward, *Slic (simple line interface calculation)*, 5th International Conference on Numerical Methods in Fluid Dynamics (A. I. van der Vooren and P. J. Zandbergen, eds.), Springer-Verlag, 1976, pp. 330–340.
- [19] W. J. Rider and D. B. Kothe, *Reconstructing volume tracking*, Journal of Computational Physics **141** (1998), 112–152.
- [20] R. Scardovelli and S. Zaleski, *Direct numerical simulation of free-surface and interfacial flow*, Annual Review of Fluid Mechanics **31** (1999), 567–603.
- [21] B. K. Swartz, *The second-order sharpening of blurred smooth borders*, Mathematics of Computation **52** (1989), no. 186, 675–714.
- [22] D. L. Youngs, *Time dependent multi-material flow with large fluid distortion*, Numerical Methods for Fluid Dynamics (K. W. Morton and M. J. Baines, eds.), Academic Press, 1982, pp. 273–285.
- [23] ———, *An interface tracking method for a 3d Eulerian hydrodynamics code*, Tech. Report AWE/44/92/35, AWRE Design and Math Division, 1984.